

SYSTEM AND METHOD FOR INFINIBAND RECEIVE FLOW CONTROL WITH COMBINED BUFFERING OF VIRTUAL LANES AND QUEUE PAIRS

5

ABSTRACT

A system and method for implementing flow control, at the link and/or transport layers, for InfiniBand receive traffic. A shared memory structure may be used for combined queuing of virtual lane and queue pair traffic. Each virtual lane is allocated a number of memory buffers; a packet is dropped if buffering it would cause its virtual lane to exceed its allocation of buffers. For each active queue pair, a linked list of buffers is maintained in the structure. Each queue pair is dedicated zero or more of its virtual lane's buffers, and may also use a set of buffers shared among multiple queue pairs. Thresholds are established in a queue pair's dedicated set of buffers and/or the shared set of buffers. As each threshold number of buffers is used, a queue pair can advertise fewer message credits. RNR-NAKs are issued when no more buffers are available to a queue pair.

10

15